REPORT DOCUMENTATION PAGE

AD-A260 783

estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503

REPORT DATE	3. REPORT TYPE AND DATES COVERED
	FINAL 15 Apr 87 – 14 Oct 92

4. TITLE AND SUBTITLE	5. FUNDING NUMBERS
"PARALLEL MATRIX COMPUTATIONS" (U)	61102F
	2304/A3

6. AUTHOR(S)
Drs G.W. Stewart & Dianne P. O'Leary

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
University of Maryland Department of Computer Science College Park MD 20742	AFOSR-TR- 93 0014

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFOSR/NM 110 Duncan Ave Suite B115 Bolling AFB DC 20332-0001	AFOSR-87-0188

DTIC ELECTE FEB 3 1993 S D C

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION AVAILABILITY STATEMENT	12b. DISTRIBUTION CODE
Approved for public release; Distribution unlimited	UL

13. ABSTRACT (Maximum 200 words)

This project concerns the design and analysis of algorithms for matrix computations to be run in a processor-rich environment. The researchers focused on algorithms that require no global control and that can be run on systems with only local connections among processors. They investigated these algorithms both theoretically and practically. The final year's work has focused on subspace updating, granularity of communication, projection methods for eigenvalue problems, polynomial preconditioners for conjugate gradients, and other topics.

93-01923

93

14. SUBJECT TERMS	15. NUMBER OF
	20
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	SAR

FINAL REPORT
AFOSR 87-0188
PARALLEL MATRIX COMPUTATIONS
1988 – 1992

D. P. O'LEARY
G. W. STEWART

Research Accomplishments   DTIC QUALITY INSPECTED 3

In this work we have pursued the development of parallel algorithms for matrix computations. We highlight a few of these activities below, and include a complete list of the publications resulting from this research grant.

Bounds on Scaled Projectors. Scaled projection operators arise in computations for weighted least squares problems, linear programming algorithms, and other applications. Let X be a matrix of full column rank and let D be a diagonal matrix with positive diagonal elements. We have shown that the weighted pseudo-inverse defined by $X_D^\dagger = (X^T D X)^{-1} X^T D$ and the associated oblique projection $P_D = X X_D^\dagger$ r, have norms bounded by numbers that are independent of D.

Iteratively Reweighted Least Squares. We studied various algorithms to obtain estimates of solution vectors and residual vectors for the linear model $Ax = b + e = b_{true}$ using an iteratively reweighted least squares criterion, which tends to diminish the influence of outliers compared with the standard least squares criterion. Algorithms appropriate for dense and sparse matrices were developed. Solving Newton's linear system using updated matrix factorizations or the (unpreconditioned) conjugate gradient iteration gave the most effective algorithms. Four weighting functions were compared, and results were obtained for sparse well-conditioned and ill-conditioned problems.

Subspace updating. An important problem in array signal processing is to determine the null space of a matrix of signals sampled at discrete times by an array of m sensors. It is necessary that the subspace be updated in real time. The customary approach has been through the singular value decomposition; however, this decomposition cannot be updated exactly in fewer than $O(m^3)$ operations, and all parallel algorithms proposed for approximate updating require $O(m^2)$ processors. Recently we have introduced a new decomposition — the URV decomposition — which can be updated sequentially in $O(m^2)$ operations, and in parallel in $O(m)$ operations using $O(m)$ processors. This decomposition has applications in other areas of signal processing. We are investigating a sequential implementation of the method. A parallel implementation has been debugged on a simulator, we are preparing to move it to the IWARP.

Parallel QR factorization. A project on parallel QR factorization has been completed. A parallel Gram-Schmidt algorithm and a parallel Householder algorithm have

been developed and programmed, and analytical models for the time complexity of these algorithms have been developed. The models were validated over a wide range of parameter values for floating point and communication speed through experiments on the ZMOB, MCMOB, a 16 processor Butterfly with hardware floating point, and a 128 processor Butterfly with software floating point.

Interprocessor communication. It has long been known that there is a close relationship between granularity of communication and granularity of computation. Roughly speaking the coarser the granularity of communication, the coarser the granularity of computation must be to compensate. In a paper to appear in *Parallel Computing*, we investigate this phenomenon theoretically and empirically. The conclusion is that to solve the large problems that tomorrow's generation of parallel computers can hold, we must have fine grained communication.

Polynomial preconditioners for conjugate gradient algorithms. Preconditioning to produce a more favorable distribution of eigenvalues is essential in using the conjugate gradient algorithm to solve linear systems of equations. The choice of preconditioner must be well matched to the problem and to the computer architecture. Polynomial preconditioning is a useful tool in the effective use of the conjugate gradient algorithm on special architectures such as message passing parallel computers, machines with hierarchical memory, vector processors, and machines with very limited memory. We have developed a new adaptive algorithm that uses a polynomial based on the residual polynomial from k steps of the conjugate gradient algorithm. This preconditioning requires no prior information about the matrix and is efficient on a variety of architectures.

Eigenvalues of Arrowhead Matrices. A query from a physicist led us to consider the problem of finding the eigenvalues and eigenvectors of a symmetric matrix with nonzeroes only in the main diagonal and the last row and column. A highly parallel algorithm was developed and an error analysis was completed.

Projection methods for eigenvalue problems. Projection methods, such as Kaczmarz's algorithm are promising for very large eigenvalue problems, since they use comparatively little storage and access the matrix only one row at a time. Kaczmarz's method for inhomogeneous systems has been extensively analyzed. In spite of this, very little is known about the rate of convergence of the method. Using a relation between Kaczmarz's method and SOR, we conjecture that the convergence rate should be approximately

$$\frac{1 - \sigma_{n-1}^2}{1 + \sigma_{n-1}^2},$$

where σ_{n-1} is the second smallest singular value of $A - \lambda I$ and λ is the eigenvalue whose eigenvector is to be found.

Markov chains. Although the theory of Gaussian elimination for the solution of Markov chains insures that it will perform well on chains with balanced transition probabilities, the method fails for nearly uncoupled Markov chains. We have given a formal rounding-error analysis of a variant of Gaussian elimination which works for these chains.

Two students completed Ph.D. theses in the study of Markov chains. Xiaobai Sun presented a unified framework for studying various aggregation algorithms and used it to study the convergence rate of these algorithms. Pil Park studied iterative algorithms for overflow queuing networks, and achieved good results with a combination of projection and preconditioning.

Perturbation theory. Six problems in matrix perturbation theory have emerged from our work. They concern the condition of nearly uncoupled Markov chains, the computation of residual bounds for eigenvalues, the computation of residual bounds for singular values, the condition of multiple eigenvalues, the perturbations of nearly transient Markov chains, and the perturbation of matrix factorizations. Satisfactory solutions of these problems have been found.

Constrained matrix Sylvester equations. The matrix Sylvester problem of finding a matrix T to satisfy $AT + TF = C$ arises in applications in control problems involving the solution of ordinary differential equations. In the design of reduced order observers for loop transfer recovery, the solution T is further constrained to satisfy $TB = 0$, but the matrix C is the product of two matrices, one to be determined. Questions of existence and uniqueness of solutions to such problems have been studied, and a computational algorithm has been developed and tested, and applied to design of reduced order observers that achieve loop transfer recovery in aircraft.

Publication, etc.

Technical reports

1. G. W. Stewart, "An Iterative Method for Solving Linear Inequalities," CS-TR-1833, University of Maryland, April, 1987

   This paper describes a method for solving homogeneous linear inequalities. The numerical techniques required by the algorithm can be parallelized and are important in a number of other applications.

2. Gene H. Golub and Dianne P. O'Leary, "Some History of the Conjugate Gradient and Lanczos Algorithms: 1948-1976," CS-TR-1859 and UMIACS-TR-87-20, University of Maryland, June, 1987.

   This manuscript gives some of the history of the conjugate gradient and Lanczos algorithms and an annotated bibliography for the period 1948-1976.

3. Robert A. van de Geijn, "Implementing the QR-Algorithm on an Array of Processors," CS-TR-1897, University of Maryland, August, 1987.

QR-algorithms for solving the algebraic eigenvalue problem that initially reduce the matrix to upper Hessenberg form and utilize traditional shifting strategies do not lend themselves to efficient implementation on a grid of processors. In this thesis, we introduce a variation of the QR-algorithm that works with the full matrix and show how it can be implemented on a square array of processors. By using a deferred shifting scheme, iterations can be pipelined, thereby reducing processor idle time. A thorough analysis of deferred shifting techniques show that the asymptotic convergence rate remains acceptable.

4. Zhaojun Bai, "The Direct GSVD Algorithm and Its Parallel Implementation," CS-TR-1901, University of Maryland, August, 1987.

The generalized singular value decomposition (GSVD) is the simultaneous reduction of any two matrices having the same number of columns to diagonal matrices by premultiplying by two different orthogonal matrices and postmultiplying by the same nonsingular matrix. Following the work of C. C. Paige on the sequential Jacobi-like GSVD algorithm, we first provide a clearer description of his algorithm and a more straightforward proof. A new version of the direct GSVD algorithm, called the direct GSVD algorithm, is given. The error analysis shows that the algorithm is stable in the presence of rounding errors.

A parallel implementation of the direct GSVD algorithm is proposed in the second part of the paper. The parallel algorithm falls into two parts. The first is that the input matrices are preprocessed in parallel by computing their upper trapezoidal forms, for which we develop a parallel QR decomposition with column pivoting. The second is the parallel computations of the GSVD of two upper trapezoidal matrices.

Finally, the direct GSVD algorithm is used to derive an efficient method for the problem of equality-constrained least squares, and show how effective the GSVD is in dealing with a linearly constrained Gauss-Markov model. The GSVD provides an efficient algorithm and reveals the structure of the model more clearly than the usual general inverse expressions.

5. Zhaojun Bai "Data-flow Algorithm for Computing the Singular Value Decomposition," CS-TR-1941 University of Maryland, August, 1987.

In this report, a data-flow algorithm for computing the singular value decomposition (SVD) of any matrix is developed. It contains the description of a preprocession QRD algorithm for QR decomposition and SVD algorithm for the SVD of a triangular matrix. Asymptotic convergence rate in parallel odd-even ordering

is examined. Moreover, the computational network, assignment of the computational nodes, and data communication complexity are also presented.

6. Zhaojun Bai, "Numerical Treatment of Restricted Gauss-Markov Model," CS-TR-1942 University of Maryland, August, 1987.

The singular value decomposition has been widely used in the ordinary linear model and other statistical problems. In this paper, we shall introduce the generalized singular value decomposition of any two matrices X and H having the same number of columns to treat large scale restricted Gauss-Markov models $(y, X beta = r, sigma^2 I)$. This method leads to an efficient, numerically stable and easily programmed algorithm for the best linear unbiased estimator in large scale restricted Gauss-Markov linear models. The added information and sensitivity that it provides may be very useful in understanding the problem.

7. G. W. Stewart, "The Crab: a Dialogue," University of Maryland CS-TR-2025, May 1988.

The CRAB is a module for building memory connected parallel systems in which processors that are connected to one another can read and write each other's memory. This arrangement solves some of the communication problems associated with more conventional message passing systems. In particular, it is possible for a line of processors to pipeline data, altering the data items as they pass by. This report contains a dialogue about the CRAB, describing what it is and what it can do.

8. G. W. Stewart, "On Scaled Projections and Pseudo-Inverses," CS-TR 2026, May, 1988.

Let X be a matrix of full column rank and let D be a diagonal matrix with positive diagonal elements. The weighted pseudo-inverse defined by $X_L^\dagger = (X^T D X)^{-1} X^T D$ and the associated oblique projection $P_D = X X_D^\dagger$ arise in many applications. In this paper, we show that the norms of both matrices are bounded by numbers that are independent of D.

9. Dianne P. O'Leary and Peter Whitman, "Parallel QR Factorization by Householder and Modified Gram-Schmidt Algorithms," CS-TR 2119, UMIACS-TR-88-78, October, 1988.

In this paper, the parallel implementation of two algorithms for forming the QR factorization of a matrix is studied. We propose parallel algorithms for the modified Gram-Schmidt and the Householder algorithms on message passing systems in which the matrix is distributed by blocks of rows. The models that predict

performance of the algorithms are validated by experimental results on several parallel machines.

10. G. W. Stewart, "Stochastic Perturbation Theory," CS-TR 2129, UMIACS-TR-88-76, October, 1988.

    In this paper, classical matrix perturbation theory is approached from a probabilistic point of view. The perturbed quantity is approximated by a first order perturbation expansion, in which the perturbation is assumed to be random. This permits the computation of statistics estimating the variation in the perturbed quantity. Up to the higher order terms that are ignored in the expansion, these statistics tend to be more realistic than perturbation bounds obtained in terms of norms. The technique is applied to a number of problems in matrix perturbation theory, including least squares and the eigenvalue problem.

11. G. W. Stewart, "Communication and Matrix Computations on Large Message Passing Systems," CS-TR-2135, University of Maryland, October, 1988.

    This paper is concerned with the consequences for matrix computations of having a rather large number of general purpose processors, say ten or twenty thousand, connected in a network in such a way that a processor can communicate only with its immediate neighbors. Certain communication tasks associated with most matrix algorithms are defined and formulas developed for the time required to perform them under several communication regimes. The results are compared with the times for a nominal n^3 floating point operations. The results suggest that it is possible to use a large number of processors to solve matrix problems at a relatively fine granularity.

12. D. P. O'Leary and G. W. Stewart, "Computing the Eigenvalues and Eigenvectors of Arrowhead Matrices," CS-TR-2203, University of Maryland, February 1989.

    This paper treats the eigenvalue problem for a symmetric matrix which is zero except for its main diagonal and one row and column. Such problems arise in the description of radiationless transitions in isolated molecules and of oscillators vibrationally coupled with a Fermi liquid. In these applications the order n of the matrix A can be in the thousands. The purpose of this paper is to present formulas and efficient and highly parallel algorithms for computing eigenvalues and eigenvectors of such matrices.

13. Dianne P. O'Leary, "On Bounds for Scaled Projections and Pseudo-Inverses," UMIACS-TR-89-32 CS-TR-2215, University of Maryland February, 1989.

    Let X be a matrix of full column rank and let D be a positive definite diagonal matrix. In a recent paper, Stewart considered the weighted pseudo-inverse $X_D^\dagger =$

$(X^TDX)^{-1}X^TD$ and the associated oblique projection $P_D = XX_D^\dagger$, and gave bounds, independent of D, for the norms of these matrices. In this note, we answer a question he raised by showing that the bounds are computable.

14. Dianne P. O'Leary, "On Iteratively Reweighted Linear Least Squares Problems," CS-TR-2278, UMIACS 89-66, University of Maryland, July, 1989.

    Several variants of Newton's method are used to obtain estimates of solution vectors and residual vectors for the linear model $Ax = b+e = b_{true}$ using an iteratively reweighted lease squares criterion, which tends to diminish the influence of outliers compared with the standard least squares criterion. Algorithms appropriate for dense and sparse matrices are presented. Solving Newton's linear system using updated matrix factorizations or the (unpreconditioned) conjugate gradient iteration gives the most effective algorithms. Four weighting functions are compared, and results are given for sparse well-conditioned and ill-conditioned problems.

15. G. W. Stewart, "Perturbation Theory and Least Squares with Errors in the Variables," UMIACS-TR-89-97, CS-TR 2326, University of Maryland, October, 1989.

    In this note we examine what matrix perturbation theory has to say about ordinary least squares estimation when the regression matrix is contaminated by random errors. The conclusion is that there is a regime in which the errors can have important, even overwhelming effects yet do not affect the validity of ordinary least squares procedures. The boundary of this regime is indicated by a diagnostic number which can be calculated from the data.

16. G. W. Stewart, "Two Simple Residual Bounds for Eigenvalues of Hermitian Matrices," UMIACS-TR-89-123, CS-TR 2364, December, 1989.

    Let A be Hermitian and let the orthonormal columns of X span an approximate invariant subspace of X. Then the residual $R = AX - XM$ ($M = X^HX$) will be small. The theorems of this paper bound the distance of the spectrum of M from the spectrum of A in terms of appropriate norms of R.

17. G. W. Stewart, "On the Sensitivity of Nearly Uncoupled Markov Chains," UMIACS-TR-90-18, CS-TR 2406, February, 1990.

    Nearly uncoupled Markov chains (aka nearly completely decomposable Markov chains) arise in a variety of applications, where they model loosely coupled systems. In such systems it may be difficult to determine the transitions probabilities with high accuracy. This paper investigates the sensitivity of the limiting distribution of the chain to perturbations in the transition probabilities. The conclusion is that nearly uncoupled Markov chains are quite sensitive to such perturbations but the perturbation of the limiting distribution is not arbitrary.

18. G. W. Stewart and G. Zhang, "Eigenvalues of Graded Matrices and the Condition Numbers of a Multiple Eigenvalue," UMIACS TR 90-31, CS TR 2420, February 1990.

This paper concerns two closely related topics: the behavior of the eigenvalues of graded matrices and the perturbation of a nondefective multiple eigenvalue. We will show that the eigenvalues of a graded matrix tend to share the graded structure of the matrix and give precise conditions insuring that this tendency is realized. These results are then applied to show that the secants of the canonical angles between the left and right invariant subspaces of a multiple eigenvalue tend to characterize its behavior when its matrix is slightly perturbed.

19. Dianne P. O'Leary, "Yet Another Polynomial Preconditioner for the Conjugate Gradient Algorithm," CS-TR-2425, UMIACS 90-36, University of Maryland, March, 1990.

Polynomial preconditioning is a useful tool in the effective use of the conjugate gradient algorithm on special architectures such as mess⁻ ,sing parallel computers, machines with hierarchical memory, vector pr ₅ rs, and machines with very limited memory. In this work we investigate the use of a new adaptive algorithm which uses the polynomial preconditioner based on the residual polynomial from k steps of the conjugate gradient algorithm.

20. G. W. Stewart, "An Updating Algorithm for Subspace Tracking," UMIACS-TR-90-86, CS-TR 2494, July 1990.

In certain signal processing applications it is required to compute the null space of a matrix whose rows are samples of a signal. The usual tool for doing this is the singular value decomposition. However, the singular value decomposition has the drawback that it requires $O(p^3)$ operations to recompute when a new sample arrives. In this paper, we show that a different decomposition, called the URV, decomposition is equally effective in exhibiting the null space and can be updated in $O(p^2)$ time. The updating technique can be run on a linear array of p processors in $O(p)$ time.

21. G. W. Stewart and G. Zhang, "On a Direct Method for the Solution of Nearly Uncoupled Markov Chains," UMIACS TR 90-95, CS TR 2504, July 1990.

This note is concerned with the accuracy of the solution of nearly uncoupled Markov chains by a direct method based on the LU decomposition. It is shown that plain Gaussian elimination may fail in the presence of rounding errors. A modification of Gaussian elimination with diagonal pivoting as well as corrections of small pivots by sums of off-diagonal elements in the pivoting columns is proposed

and analyzed. It is shown that the accuracy of the solution is affected by two condition numbers associate with the aggregate and the coupling respectively.

22. G. W. Stewart, "Perturbation Theory for the Singular Value Decomposition," UMIACS-TR-90-124, CS-TR 2539, September, 1990.

The singular value decomposition has a number of applications in digital signal processing. However, the the decomposition must be computed from a matrix consisting of both signal and noise. It is therefore important to be able to assess the effects of the noise on the singular values and singular vectors — a problem in classical perturbation theory. In this paper we survey the perturbation theory of the singular value decomposition.

23. Jewel B. Barlow, Moghen M. Monahemi, and Dianne P. O'Leary, "Constrained matrix Liapunov equations," UMIACS-91-16, CS-2599, January 1991.

We consider the problem of finding matrices L and T satisfying $TA - FT = LC$ and $TB = 0$. We establish existence conditions for the solution, derive an algorithm for computing the solution, and discuss conditions under which the matrix $[C^T, T^T]$ is full rank. The problem arises in control theory, in the design of reduced order observers that achieve loop transfer recovery.

24. Jewel B. Barlow, Moghen M. Monahemi, and Dianne P. O'Leary, "The design of reduced-order Luenberger observers with precise LTR," UMIACS-91-17, CS-2600, January 1991.

This work concerns the design of reduced-order observers for controllable, observable, and regular systems in which the number of measurements is more than the number of controls. It uses eigenstructure assignment whereas other approaches use Kalman filter (LQG/LTR) methods. The advantages of this approach are precise rather than approximate LTR, no restriction to minimum phase systems, finite rather than infinite observer gain, and simpler and more efficient numerical calculation. Case studies are presented illustrating these features.

25. G. W. Stewart "On an Algorithm for Refining a Rank-Revealing URV Factorization and a Perturbation Theorem for Singular Values" UMIACS-TR-91-38, CS-TR 2626, March 1991

In this note we consider an iterative algorithm for moving a triangular matrix toward diagonality. The method is shown to converge under conditions that are likely to be met in practice. A result of the convergence proof in a new perturbation theorem for singular values.

26. G. W. Stewart "Updating A Rank-Revealing ULV Decomposition" UMIACS-TR-91-39, CS-TR 2627, March 1991

A ULV decomposition of a matrix A of order n is a decomposition of the form $A = ULV^H$, where U and V are orthogonal matrices and L is a lower triangular matrix. When A is approximately of rank k, the decomposition is rank revealing if the last $n - k$ rows of L are small. This paper presents algorithms for updating a rank-revealing ULV decomposition. The algorithms run in $O(n^2)$ time, and can be implemented on a linear array of processors to run in $O(n)$ time

27. G. Adams, M. F. Griffin, and G. W. Stewart "Direction-of-Arrival Estimation Using the Rank-Revealing URV Decomposition" UMIACS-TR 91-46, CS-TR-2640, March 1991

An algorithm for updating the null space of a matrix is described. The algorithm is based on a new decomposition, called the URV decomposition, which can be updated in $O(N^2)$ and serves as an intermediary between the QR decomposition and the singular value decomposition. The URV decomposition is applied to a high-resolution direction of arrival problem based on the MUSIC algorithm. A virtue of the updating algorithm is the running estimate of rank.

28. G. W. Stewart "Error Analysis of QR Updating with Exponential Windowing" UMIACS-TR-91-79, CS-TR 2685, May 1991

Exponential windowing is a widely used technique for suppressing the effects of old data as new data is added to a matrix. Specifically, given an $n \times p$ matrix X_n and a "forgetting factor" $\beta \in (0, 1)$, one works with the matrix $\mathrm{diag}(\beta^{n-1}, \beta^{n-2}, \ldots, 1)X_n$. In this paper we examine an updating algorithm for computing the QR factorization of $\mathrm{diag}(\beta^{n-1}, \beta^{n-2}, \ldots, 1)X_n$ and show that it is unconditionally stable in the presence of rounding errors.

29. Pill Seong Park, "Iterative Solution of Sparse Singular Systems of Equations Arising from Queuing Networks," UMIACS-TR-91-83, CS-TR-2690, June 1991 ·

Iterative methods for solving large sparse singular systems $Ax = 0$ arising from queuing networks having overflow capacity are presented. For such overflow models, no analytic solution exists and the Kolmogorov balance equation has to be solved explicitly. The resulting matrix is irreducible, non-symmetric, and has a one dimensional null space. However, the matrix is sparse, highly structured, and has property-A.

We transform the problem into an eigenvalue problem $Tx = x$, where the eigenvector corresponding to the eigenvalue 1 is the desired solution. The choice of the Jacobi iteration leads to a 2-cyclic algorithm which reduces the necessary amount of work by 1/2, and computation of the residual needs no extra work.

Inspired by the similarity between Markov models of queuing networks and the grid problems arising from discretization of partial differential equations, a few

aggregation/disaggregation(A/D) type methods with some ideas from geometric multigrid methods are considered.

More effective methods to accelerate convergence of the 2-cyclic algorithm are discussed. The methods employ orthogonal projectors onto the subspace spanned by dominant eigencomponents in the residual. The projection step is further refined by Arnoldi's method without extra matrix-vector multiplication by using the results of power iterations. Adopting the Chebyshev iterations as a main driving force combined with power iterations and projection steps refined by Arnoldi's method, the resulting hybrid algorithm outperforms the Chebyshev iteration methods with optimal parameters. We study the convergence of the hybrid algorithm and look for conditions when the projection step can accelerate convergence of the underlying method. Numerical experiments provide further evidence that the methods can be quite efficient, especially for harder problems.

30. G. W. Stewart "Perturbation Theory for Rectangular Matrix Pencils" UMIACS-TR-91-105, CS-TR 2721, July 1991

The theory of eigenvalues and eigenvectors of rectangular matrix pencils is complicated by the fact that arbitrarily small perturbations of the pencil can cause them disappear. However, there are applications in which the properties of the pencil ensure the existence of eigenvalues and eigenvectors. In this paper it is shown how to develop a perturbation theory for such pencils.

31. Per Christian Hansen and Dianne Prost O'Leary, "The use of the L-curve in the regularization of discrete ill-posed problems," UMIACS-91-142, CS-2781, October 1991.

In order to produce reasonable solutions to ill-posed problems, regularization algorithms are often used. The L-curve is a plot—for all valid regularization parameters—of the size of the regularize d solution versus the size of the corresponding residual. We establish two main results. First we give a unifying charac terization of various regularization methods and show that the measurement of "size" is dependent on the particular regularization method chosen; for example, the 2-norm is appropriate for Tikhonov regularization, but a 1-norm in the coordinate system of the singular value decomposition (SVD) is relevant to truncated SVD regularization . Secondly, we propose a new method for choosing the regularization parameter based on the L-curve, and show how this method can be implemented efficiently. We compare the method to generalized cross validation and demonstrate that our new method is more robust in the presence of correlated errors.

32. Chiou-Ming Huang and Dianne P. O'Leary, A Krylov Multisplitting Algorithm for Solving Linear Systems of Equations Inst. for Mathematics and Its Applications,

University of Minnesota Preprint 992, July 1992.

We consider the practical implementation of Krylov subspace methods (conjugate gradients, GMRES, etc.) for parallel computers in the case where the preconditioning matrix arises from a multisplitting. We show that the algorithm can be efficiently implemented by dividing the work into tasks that generate search directions and a single task that minimizes over the resulting subspace. Each task is assigned to a subset of processors. It is not necessary for the minimization task to send frequent information to the direction generating tasks, and this leads to high utilization with a minimum of synchronization. We study the convergence properties of various forms of the algorithm and present results of numerical examples on a sequential computer. We consider the practical implementation of Krylov subspace methods (conjugate gradients, GMRES, etc.) for parallel computers in the case where the preconditioning matrix arises from a multisplitting. We show that the algorithm can be efficiently implemented by dividing the work into tasks that generate search directions and a single task that minimizes over the resulting subspace. Each task is assigned to a subset of processors. It is not necessary for the minimization task to send frequent information to the direction generating tasks, and this leads to high utilization with a minimum of synchronization. We study the convergence properties of various forms of the algorithm and present results of numerical examples on a sequential computer.

33. J. Barlow, M. Monahemi, and D. P. O'Leary, Constrained Matrix Sylvester Equations Computer Science Department Report CS-?, Institute for Advanced Computer Studies Report UMIACS-91-?, University of Maryland, 1991. We consider the problem of finding matrices L and T satisfying $TA - FT = LC$ and $TB = 0$. We establish existence conditions for the solution, derive an algorithm for computing the solution, and discuss conditions under which the matrix $[C^T, T^T]$ is full rank. The problem arises in control theory, in the design of reduced order observers that achieve loop transfer recovery.

34. J. Barlow, M. Monahemi, and D. P. O'Leary, The Design of Reduced-Order Observers with Precise Loop Transfer Recovery, Computer Science Department Report CS-?, Institute for Advanced Computer Studies Report UMIACS-92-?, University of Maryland, 1991 This paper concerns the design of reduced-order observers for systems in which the number of measurements is more than the number of controls. We develop an algorithm that applies to regular systems that have no transmission zeroes. The algorithm uses eigenstructure assignment whereas other approaches use Kalman filter methods. The advantages of this approach are the following: i) precise loop transfer recovery rather than approximate loop transfer recovery, ii) finite observer gain rather than asymptotic observer gain, iii) modest

computational tools and operations counts. Case studies are presented illustrating these features.

35. Dianne P. O'Leary, Iterative Methods for Finding the Stationary Vector for Markov Chains, Institute for Mathematics and Its Applications, University of Minnesota, Preprint 932, February 1992

This overview concerns methods for estimating the steady-state vector of ergodic Markov chains. The problem can be cast as an ordinary eigenvalue problem, but since the eigenvalue is known, it can equally well be studied as a nullspace problem or as a linear system. We discuss iterative methods for each of these three formulations. Many of the applications, such as queuing modeling, have special structure that can be exploited computationally, and we give special emphasis to three ideas for exploiting this structure: decomposibility, separability, and multi-level aggregation. Such ideas result in a large number of diverse algorithms, many of which are poorly understood.

36. K.J.R. Liu, G.W. Stewart, and Y.-J. J. Wu, URV ESPRIT for Tracking Time-Varying Signals Computer Science Department Report CS-?, Institute for Advanced Computer Studies Report UMIACS-92-?, University of Maryland, October 1992.

ESPRIT is an algorithm for determining the fixed directions of arrival of a set of narrowband signals at an array of sensors. Unfortunately, its computational burden makes it unsuitable for real time processing of signals with time-varying directions of arrival. In this work we develop a new implementation of ESPRIT that has potential for real time processing. It is based on a rank-revealing URV decomposition, rather than the eigendecomposition or singular value decomposition used in previous ESPRIT algorithms. We demonstrate its performance on simulated data representing both constant and time-varying signals. We find that the URV-based ESPRIT algorithm is effective for estimating time-varying directions-of-arrival at considerable computational savings over the SVD-based algorithm.

37. G. W. Stewart, On the Perturbation of Markov Chains with Nearly Transient States, UMIACS-TR-92-14, CS-TR-2835, January 1992.

Let A be an irreducible stochastic matrix of the form

$$A = \begin{pmatrix} A_{11} & E_{12} \\ A_{21} & A_{22} \end{pmatrix}.$$

If E_{22} were zero, the states corresponding to A_{22} would be transient in the sense that if the steady state vector y^T is partitioned conformally in the form $(y_1^T \ y_2^T)$

then $y_2^T = 0$. If F_{22} is small, then y_2^T will be small, and the states are said to be nearly transient. It this paper it is shown that small relative perturbations in A_{11}, A_{21}, and A_{22}, though potentially larger than y_2^T, induce only small relative perturbations in y_2^T.

38. G. W. Stewart, On the Perturbation of LU, Cholesky, and QR Factorizations, UMIACS-TR-92-24, CS-TR-2848, February 1992

In this paper error bounds are derived for a first order expansion of the LU factorization of a perturbation of the identity. The results are applied to obtain perturbation expansions of the LU, Cholesky, and QR factorizations.

39. G. W. Stewart, Updating URV Decompositions in Parallel, UMIACS-TR-92-44, CS-TR-2880, April 1992

A URV decomposition of a matrix is a factorization of the matrix into the product of a unitary matrix (U), an upper triangular matrix (R), and another unitary matrix (V). In an earlier paper [UMIACS-TR-90-86] it was shown how to update a URV decomposition in such a way that it reveals the effective rank of the matrix. It was also argued that the updating procedure could be implemented in parallel on a linear array of processors; however, no specific algorithms were given. This paper gives a detailed implementation of the updating procedure.

40. Z. Bai and G. W. Stewart, SRRIT — A FORTRAN Subroutine to Calculate the Dominant Invariant, Subspace of a Nonsymmetric Matrix, UMIACS TR-92-61, CS TR-2908, May, 1992

SRRIT is a FORTRAN program to calculate an approximate orthonormal basis for a dominant invariant subspace of a real matrix A by the method of simultaneous iteration Specifically, given an integer m, SRRIT attempts to compute a matrix Q with m orthonormal columns and real quasi-triangular matrix T of order m such that the equation

$$AQ = QT$$

is satisfied up to a tolerance specified by the user. The eigenvalues of T are approximations to the m largest eigenvalues of A, and the columns of Q span the invariant subspace corresponding to those eigenvalues. SRRIT references A only through a user provided subroutine to form the product AQ; hence it is suitable for large sparse problems.

41. G. W. Stewart, Determining Rank in the Presence of Error, UMIACS TR-92-108, CS TR-2972, October 1992

The problem of determining rank in the presence of error occurs in a number

of applications. The usual approach is to compute a rank-revealing decomposition and make a decision about the rank by examining the small elements of the decomposition. In this paper we look at three commonly use decompositions: the singular value decomposition, the pivoted QR decomposition, and the URV decomposition.

Presentations

1. G. W. Stewart, "Parallel Matrix Computations," A series of lectures presented at Fudan University, Shanghai; Nanjing University, Nanjing; and the Chinese Academy of Science, Beijing, January-February, 1987.

2. D. P. O'Leary, "Some Small Problems in Parallel Computing," University of Illinois, January, 1987.

3. D. P. O'Leary, "Some Small Problems in Parallel Computing," Pennsylvania State University, March, 1987.

4. G. W. Stewart, "Communications and Matrix Computations," Workshop on Numerical Computation, Northwestern University, March, 1987.; Argonne National Laboratories, March, 1987; University of Illinois, March, 1987.

5. D. P. O'Leary, "Domino: A Portable Parallel Programming Environment," Supercomputing Research Center, March, 1987.

6. G. W. Stewart, "Jim Wilkinson's Contributions to Rounding-Error Analysis," (invited) SIAM International Conference, Paris, July, 1987.

7. G. W. Stewart, "A Method for Solving Linear Inequalities," (invited), NPL Conference in memory of James H. Wilkinson, Teddington, July, 1987.

8. D. P. O'Leary, "Some Small Problems in Parallel Processing," (invited) Gatlinburg X Meeting on Numerical Linear Algebra, Fairfield Glade, Tennessee, October, 1987.

9. G. W. Stewart, "Communications and Matrix Computations on a Rather Large Number of Processors," (invited) Joint Statistical Meeting, San Francisco, August, 1987. Also presented at the National Bureau of Standards, September, 1987, and Los Alamos National Laboratory October, 1987.

10. G. W. Stewart "Parallel Linear Algebra in Statistical Computation," (invited) COMPSTAT conference, Copenhagen, August 1988.

11. G. W. Stewart "Distributed Basic Linear Algebra Subprograms," Supercomputing Research Center, October, 1988. Also at the University of Texas, and Rice University.

12. D. P. O'Leary, "On Iteratively Reweighted Least Squares Problems, " SIAM Sparse Matrix Symposium, Gleneden Beach, Oregon, May, 1989.

13. G. W. Stewart, "Perturbation Theory and Least Squares with Errors in the Variables," AMS Conference on Measurement Error Models, Humboldt, June, 1989.

14. G. W. Stewart, "On the Sensitivity of Nearly Uncoupled Markov Chains," First International Workshop on the Numerical Solution of Markov Chains, Raleigh, NC, January, 1990.

15. G. W. Stewart, "Perturbation Theory for the Singular Value Decomposition," Second International Conference on the Singular Vaule Decomposition in Signal Processing, Kingston, RI, 1990.

16. G. Adams, M. F. Griffin, and G. W. Stewart, "Direction-of-Arrival Estimation Using the Rank-Revealing URV Decomposition," ICAASP91, Toronto, April, 1991.

17. D. P. O'Leary, "Regularizing Iterations for the Solution of Discrete Ill-Posed Problems," Numerical Linear Algebra Conference, Mathematisches Forschungsinstitut Oberwolfach, Germany, April, 1991.

18. G. W. Stewart, "Updatable, Rank Revealing Factorization," International Conference on Total Least Squares, Leuven, Belgium, August, 1991.

19. J. Barlow, M. Monahemi, and D. P. O'Leary "The Design of Reduced Order Luenberger Observers with Precise LTR," *Proceedings of the AIAA Meeting on Guidance, Navigation and Control* New Orleans, August 1991.

20. J. Barlow, M. Monahemi, and D. P. O'Leary "Considerations on Loop Transfer Recovery for Non-minimum Phase Plants," *Proceedings of the AIAA Aircraft Design Systems and Operations Meeting*, Baltimore, September 1991, AIAA-91-3086.

21. G. W. Stewart, "The Numerical Treatment of Markov Chains," SIAM Conference on Applied Linear Algebra, Minneapolis, September, 1991.

22. D. P. O'Leary, "Constrained Matrix Sylvester Equations," SIAM Conference on Applied Linear Algebra, Minneapolis, Sept ember, 1991.

23. D. P. O'Leary, "Iterative Methods for Solving Markov Chains," Workshop on Linear Algebra, Markov Chains, and Queuing Problems, Institute for Mathematics and Its Applications, University of Minnesota, Minneapolis, January 1992.

24. Chiou-Ming Huang and D. P. O'Leary, "Preconditioning Parallel Multisplittings for Solving Linear Systems of Equations," International Conference on Supercomputing (Washington, DC, July 1992) ACM Press, New York, 1992, 478-484.

25. M. Monahemi, J. Barlow, and D. P. O'Leary, "On the Precise Loop Transfer Recovery and Transmission Zeroes," First IEEE Conference on Control Applications, Dayton, Ohio, September 1992.

26. D. P. O'Leary, (invited) "Iterative Methods for Finding the Stationary Vector for Markov Chains," *Linear Algebra, Markov Chains, and Queuing Models (IMA, January 1992)*, Carl Meyer and Robert Plemmons, eds., Springer-Verlag, (IMA Volumes in Math. and Its Applics. Vol. 48) New York, 1993?, 125-???.

27. D. P. O'Leary, (invited) "Some Themes in Gene H. Golub's Work on Iterative Methods," *Iterative Methods for Sparse and Structured Problems (IMA, February 1992)*, Gene H. Golub, Anne Greenbaum, and Mitchell Luskin, eds., Springer-Verlag, New York, 1993?

28. G. W. Stewart, (invited) "Gaussian Elimination, Perturbation Theory, and Markov Chains." IMA Workshop on Matrix Theory, Probability Theory, and Markov Chains, Minneapolis, MN 1992.

29. G. W. Stewart, (invited) "A Two-Stage Iteration for Solving Nearly Uncoupled Markov Chains." IMA Workshop on Iterative Methods for Sparse and Structured Systems, Minneapolis, MN 1992.

30. G. W. Stewart, (invited) "Updatable, Rank Revealing Factorizations," SIAM Regional Conference on Numerical Analysis, Kent, OH, 1992.

31. G. W. Stewart, (invited) "Perturbation Theory for Markov Chains," Second Conference of the International Linear Algebra Society, Lisbon, Portugal, 1992.

32. G. W. Stewart, (invited) "Detecting Rank Degeneracy in the Presence of Errors," NATO Workshop on Large Scale Linear Algebra, Leuven, Belgium, 1992. :

33. G. W. Stewart, (invited) "The Numerical Treatment of Markov Chains," NATO Workshop on Large Scale Linear Algebra, Leuven, Belgium, 1992.

Publications

1. G. W. Stewart, "Computing the CS Decomposition of a Partitioned Orthonormal Matrix," *Numerische Mathematik* **40** (1982) 297-306.

2. D. P. O'Leary, "Ordering schemes for parallel processing of certain mesh problems," *SIAM Journal on Scientific and Statistical Computing* **5** (1984) 620-632.

3. D. P. O'Leary, G. W. Stewart, "Data-flow algorithms for parallel matrix computations," *Communications of the ACM* **28** (1985) 840-853.

4. G. W. Stewart, "A Jacobi-like Algorithm for Computing the Schur Decomposition of a Non-Hermitian Matrix," *SIAM Journal on Scientific and Statistical Computing* **6** (1985) 853-864.

5. D. P. O'Leary, R. E. White, "Multi-splittings of matrices and parallel solution of linear systems," *SIAM Journal on Algebraic and Discrete Methods* **6** (1985) 630-640.

6. G. W. Stewart "A Note on Complex Division," *ACM Transactions on Mathematical Software* **11** (1985) 238-241.

7. D. P. O'Leary and G. W. Stewart, "Assignment and Scheduling in Parallel Matrix Factorization," *Linear Algebra and Its Applications* **77** (1986) 275-300.

8. G. W. Stewart, "Communication in Parallel Algorithms: An Example," *Proceedings of the 18th Symposium on the Interface between Computer Science and Statistics,* American Statistical Association, *1986,* 11-14.

9. D. P. O'Leary, R. van de Geijn, G. W. Stewart, "DOMINO: A Transportable Operating System," *Parallel Processing and Medium-Scale Multiprocessors* (Proceedings of a 1986 Conference), Arthur Wouk (Ed.), SIAM Press, Philadelphia (1989) 25-34.

10. D. P. O'Leary, "Systolic Arrays for Matrix Transpose and Other Reorderings," *IEEE Transactions on Computers,* **C-36** (1987) 117-122.

11. G. W. Stewart, "A Parallel Implementation of the QR Algorithm," *Parallel Computing* **5** (1987) 187-196.

12. D. P. O'Leary, "Fine and Medium Grained Parallel Algorithms for Matrix QR Factorization," *Algorithms and Applications on Vector and Parallel Computers,* H.J.J. te Riele, Th.J. Dekker and H.A. van der Vorst, eds., Elsevier Science Publishers B.V. (North Holland), (1987) 347-349.

13. D. P. O'Leary, "Parallel Implementation of the Block Conjugate Gradient Algorithm," *Parallel Computing,* to **5** (1987) 127-140.

14. D. P. O'Leary and G. W. Stewart, "From Determinacy to Systaltic Arrays," *IEEE Transactions on Computers* **c-36** (1987) 1355-1359.

15. David C. Fisher, "Your Favorite Parallel Algorithms Might Not Be As Fast As You Think," *IEEE Transactions on Computers*, **C-37** (1988) 211-213.

16. Gene H. Golub and Dianne P. O'Leary, "Some History of the Conjugate Gradient and Lanczos Algorithms: 1948-1976," *SIAM Review*, 31 (1989) 50-102.

17. John M. Conroy, "A note on the parallel Cholesky factorization of wide banded matrices," *Parallel Computing* 10 (1989) 239-246.

18. John M. Conroy, "Parallel algorithms for the solution of narrow banded systems," *Applied Numerical Mathematics* 5 (1989) 409-421.

19. G. W. Stewart "On Scaled Projections and Pseudo Inverses," *Linear Algebra and Its Applications* **112** (1989) 189-194.

20. G. W. Stewart, G. N. Stenbakken, and T. M. Souders "Ambiguity Groups and Testability," *IEEE Transactions on Instrumentation* **38** (1989) 941-947.

21. D. P. O'Leary, "On bounds for scaled projections and pseudo-inverses," *Linear Algebra and Its Applications* 132 (1990) 115-117.

22. John M. Conroy, "Parallel Nested Dissection," *Parallel Computing* 16 (1990) 139-156.

23. D. P. O'Leary and G.W. Stewart, "Computing the eigenvalues and eigenvectors of symmetric arrowhead matrices," *Journal of Computational Physics* **90** (1990) 497-505.

24. D. P. O'Leary, "Robust regression computation using iteratively reweighted least squares," *SIAM Journal of Matrix Analysis and Applications* 11 (1990) 446-480.

25. G. W. Stewart, "Stochastic Perturbation Theory," *SIAM Review* 32 (1990) 576-610.

26. G. W. Stewart, "Communication and matrix computation on large message passing systems." *Parallel Computing* 16 (1990) 27-40

27. D. P. O'Leary and P. Whitman, "Parallel QR factorization by Householder and modified Gram-Schmidt algorithms," *Parallel Computing*, 16 (1990) 99-112.

28. G. W. Stewart, "Two simple residual bounds for the eigenvalues of Hermitian matrices," *SIAM J. Mat. Anal. Appl.* 12 (1991) 205-208.

29. D. P. O'Leary, "Yet another polynomial preconditioner for the conjugate gradient algorithm," *Linear Algebra and Its Applications*, 154 (1991) 377-388.

30. G. W. Stewart and G. Zhang, "Eigenvalues of graded matrices and the condition numbers of a multiple eigenvalue," *Numerische Mathematik* 58 (1991) 703–712.

31. G. W. Stewart and G. Zhang, "On a Direct Method for the Solution of Nearly Uncoupled Markov Chains," *Numerische Mathematik* 59 (1991) 1–12.

32. G. W. Stewart, "An Updating Algorithm for Subspace Tracking," *IEEE Trans. On Signal processing* 40 (1992) 1535–1541.

33. G. W. Stewart, "Error Analysis of QR Updating with Exponential Windowing," *Mathematics of Computation* 59 (1992) 135–140.

34. G. W. Stewart, "Updating A Rank-Revealing ULV Decomposition," *SIAM Journal on Matrix Analysis and Applications,* to appear.

35. Jewel B. Barlow, Moghen M. Monahemi, and Dianne P. O'Leary, "Constrained matrix Sylvester equations," *SIAM Journal on Matrix Analysis and Applications,* 13 (1992) 1-9.

36. Jewel B. Barlow, Moghen M. Monahemi, and D. P. O'Leary, "The design of reduced order observers with precise loop transfer recovery," *AIAA Journal of Guidance, Control, and Dynamics* to appear.

37. Chiou-Ming Huang and D. P. O'Leary, "A Krylov Multisplitting Algorithm for Solving Linear Systems of Equations," *Linear Algebra and Its Applications,* to appear.

38. P. Schweitzer and G. W. Stewart, "The Laurent Expansion of Pencils that are Singular at the Origin." To appear *Linear Algebra and Its Applications.*

39. G. W. Stewart, "On an Algorithm for Refining Rank-Revealing URV Factorization and a Perturbation Theorem for Singular Values.' To appear *Linear Algebra and Its Applications.*

40. G. W. Stewart, "Updating A Rank-Revealing ULV Decomposition." To appear *SIAM Journal on Matrix Analysis and Applications.*

41. G. W. Stewart, "On the Perturbation of LU, Cholesky, and QR Factorizations." To appear *SIAM Journal on Matrix Analysis and Applications.*

42. G. W. Stewart, "On the Perturbation of Markov Chains with Nearly Transient States." To Appear in *Numerische Mathematik.*